

Thursday, September 10, 2009

RBridges integrated to Opensolaris Build 124

RBridges? WTF? Did you ever hear of TRILL? The Transparent Interconnection of Lots of Links? Perhaps you've heard of the protocol that the TRILL WG at the IETF hopes to supersede: The Spanning Tree Protocol (STP). Almost everyone configuring networks had her or his own share of problems with this protocol and wished it into hell. But: The STP has an important purpose. Guaranteeing a loopfree logical network, even when there are loops in the physical cabling. This is important.

One of it's components is the RBridge - the routing bridge. I've heard the first time at the CEC 2007 about this concepts when Radia Perlman held a talk in the "Luminaries"-Track.

Hmm ... i know it's the oversimplification of the year, but the basic idea of TRILL and RBriges is the usage of well-proven IP routing protocols and their mechanisms to build optimal routing tables to the domain of bridging and thus getting rid of the STP. STP is an proven and venerable concept, but it has many shortfalls: slow convergence against topology changes (vulgo: cable or switch fails) especially when you use it in it's pure form, the problem of multiple paths (As STP is a tree, it forces all traffic on one line and it's not using alternate links, when you do something like that you have to work with multiple spanning trees in seperate VLANs), furthermore the traffic flows aren't optimal.

RBridges try to solve those problems. When you want to know more about this i just want to link you to the Infocom 2004 paper by Radia describing the concept, the RFC with the problem statement and the base specification of the RBridges. I just want to quote the abstract of the last document to give you an impression what the IETF wants to accomplish with this routing bridges: RBridges provide optimal pair-wise forwarding with zero configuration, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. They achieve these goals using IS-IS routing and encapsulation of traffic with a header that includes a hop count. RBridges are compatible with previous IEEE 802.1 customer bridges as well as IPv4 and IPv6 routers and end nodes. They are as invisible to current IP routers as bridges are and, like routers, they terminate the bridge spanning tree protocol.

The design supports VLANs and optimization of the distribution of multi-destination frames based on VLAN and IP derived multicast groups. It also allows forwarding tables to be sized according to the number of RBridges (rather than the number of end nodes), which allows internal forwarding tables to be substantially smaller than in conventional bridges. Why do i tell you all this stuff: Those RBridges found their way into OpenSolaris with the putback of the code from of the developments efforts from the Bridging Project (PSARC 2008/055 and PSARC 2009/344) and the RBridge Project (PSARC 2007/596). You can read more about this new developments that were integrated to Build 124 today at the webpages of the ARC cases.

Posted by Joerg Moellenkamp in English, Oracle, Solaris, Technology at 22:19

Mmh.. ohne STP wäre das Leben nur noch halb so aufregend. Zudem gibt es schon heute keinen wirklichen Grund mehr auf STP zu setzen, sofern man konsequent auf Routing setzt.

Anonymous on Sep 13 2009, 10:49